

FILE COPY

②

AD-A219 920

DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

2a. SECURITY CLASSIFICATION AUTHORITY N/A		1b. RESTRICTIVE MARKINGS N/A	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A		3. DISTRIBUTION/AVAILABILITY OF REPORT UNLIMITED	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR TR 89-032	
6a. NAME OF PERFORMING ORGANIZATION UNIVERSITY OF SOUTHERN CALIFORNIA	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION USAF, AFSC	
6c. ADDRESS (City, State, and ZIP Code) UNIVERSITY PARK LOS ANGELES, CA. 90089-1147		7b. ADDRESS (City, State, and ZIP Code) AIR FORCE OFFICE OF SCIENTIFIC RESEARCH BUILDING 410 BOLLING AFB, D.C. 20332-6448	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL (If applicable) NM	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F49620-89-C-0021	
8c. ADDRESS (City, State, and ZIP Code) Bldg 410 Bolling AFB, DC 20332-6448		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO. 61102F	PROJECT NO. 2304
		TASK NO. A7	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) RESEARCH IN KNOWLEDGE DELIVERY			
12. PERSONAL AUTHOR(S) HOVY, EDUARD H.			
13a. TYPE OF REPORT FINAL TECHNICAL	13b. TIME COVERED FROM 11/86 TO 10/89	14. DATE OF REPORT (Year, Month, Day) 1990, February 28	15. PAGE COUNT 30
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
19. ABSTRACT This report summarizes the research and development work done over four years toward the goal of automatically planning and generating fluent multisentence paragraphs of English text, while ensuring that the grammar is adequate to support the parsing of English text. The work consisted of three principal components, namely text structuring, parsing, and knowledge representation. A theory of text structure, and an accompanying text planner, were developed and successfully used to generate paragraphs in three different application domains. To ensure bidirectionality, an existing prototype parser was adapted and refined and tested on a functional grammar in to investigate the invertibility of the grammar. Knowledge representation work focused on linking the generator with arbitrary applications by developing a very general underlying taxonomy of conceptual entities which can be linked with various specific domain-related taxonomies.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL DR H WAKSMAN		22b. TELEPHONE (Include Area Code) (202) 767-5028	22c. OFFICE SYMBOL NM

RESEARCH IN KNOWLEDGE DELIVERY

Final Technical Report
AFOSR Contract F49620-87-C-0005
11/1/86 - 10/31/89

Eduard H. Hovy

University of Southern California
Information Sciences Institute

February 1990

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Code <input checked="checked" type="checkbox"/>	
Dist	Avail and/or Special
A-1	

Abstract

This report summarizes the research and development work done over four years toward the goal of automatically planning and generating fluent multisentence paragraphs of English text, while ensuring that the grammar is adequate to support the parsing of English text. The work consisted of three principal components, namely text structuring, parsing, and knowledge representation. A theory of text structure, and an accompanying text planner, were developed and successfully used to generate paragraphs in three different application domains. To ensure bidirectionality, an existing prototype parser was adapted and refined and tested on a functional grammar in to investigate the invertibility of the grammar. Knowledge representation work focused on linking the generator with arbitrary applications by developing a very general underlying taxonomy of conceptual entities which can be linked with various specific domain-related taxonomies.



1 Objectives of the Research Effort

In 1985 USC/ISI proposed a four year plan for work on knowledge delivery to conduct research to enable computers to express information in fluent multiparagraph English, while ensuring that the developed technology would also be of use for parsing. This is the final summary of the work conducted under this plan. The principal research themes are:

- Text Structure
- Parsing
- Knowledge Representations

This work was to be performed within the context of the Penman natural language generation group, which was actively constructing a natural language sentence generation program under Darpa funding.

1.1 Text Structure

The technology of text generation would remain weak and ineffective if it were restricted to single isolated sentences. Although single-sentence generation is important, it has always been simply a step toward the ability to create larger texts. As a theoretical problem, we have to understand how texts are coherently built up out of sentences, what special effects arise from using combinations of sentences, and how particular organizations of text should be selected or constructed. Until recently, text generation research has been hampered by a lack of suitable descriptive theory; existing descriptions have been too informal and too literary to be computationally useful.

The heart of the problem is that of text coherence. Coherent text can be defined as text in which the hearer knows how each part of the text relates to the whole; i.e., (a) the hearer knows why it is said, and (b) the hearer can relate the semantics of each part to a single overarching framework.

The problem of text coherence can be characterized in specific terms as follows. Assuming that the typical output of data processing procedures (such as data base searches or expert system runs) is a set of sentence- or clause-sized chunks of representation, the question is how to structure this output in a coherent multisentence paragraph. Since the permutation set of the elements defines the space of possible paragraphs, a simplistic, brute-force way to achieve coherent text would be to search this space and pick out the coherent paragraphs. Even if a well-defined criterion of coherence could be found, such a search would be factorially expensive. For example, in a paragraph of 7 input clusters, there would be $7! = 5,040$ candidate paragraphs. It is possible, however, to limit the search to a manageable size by utilizing the constraints imposed by coherence relations that hold between successive pieces of text. These relations can be formulated as search operators

and used in a hierarchical-expansion planner to limit and guide the search and to produce structures describing the coherent paragraphs. The state of this research is described in Section 2.1.

1.2 Parsing

The central task of the parsing research task was to add the functionality of language analysis (parsing) to our existing large systemic generation-oriented grammar of English. This was attempted for several reasons, the most prominent being the widespread desire among AI and Computational Linguistics workers for bidirectional grammars in which the two directions (analysis and synthesis) are compatible in both theoretical orientation and detail.

The approach taken was to re-express the grammar, Nigel, in the Functional Unification Grammar framework, then to attempt to parse by techniques which are a variant of existing unification parsing methods. At the beginning there were several issues:

1. Bidirectionality: Could the inverse of the generation grammar be found?
2. Efficiency: Unconstrained Unification has a reputation for being exponentially slow in principle and extremely slow in practice. Would this be true of an inverted Nigel?
3. Grammatical Specificity: No Systemic grammar had ever been examined for ambiguity behavior. Would we find large factors of preventable ambiguity?
4. Inversion-specific Inefficiencies: Does analysis using a grammar designed only for generation have unsuspected deficiencies in available information?

This work led to the insight that if the grammar were to be re-represented in the knowledge representation system Loom, Loom's built-in classification mechanism could be used as the central inference operation of the parser. This step is possible due to the formal relationship between the operations of classification and of subsumption (of which unification is a variant). Though this work has not been completed, this insight is very exciting and has generated much interest in the field of Computational Linguistics. The prospects are described in Section 2.2.

1.3 Knowledge Representations

As in all the symbolic disciplines, development of notation is a central part of progress in the art. Even more than many other AI applications, text generation is sensitive to the difficult and crucial problems with notation. Knowledge representation, a major subdiscipline of AI, consists almost entirely of theoretical and experimental studies of notations for information.

Although many computationally useful knowledge notations exist, their collective scope is far from comprehensive, due partly to inefficiency, and partly to formal limits. Some

notations are very general, such as those based on first- (and higher-) order predicate calculus and the lambda calculus, but are difficult to work with in practical systems; others are more convenient and efficient, but apply only to relatively narrow domains of knowledge. Many varieties of information which are representable in principle using existing notations often have no computationally tractable notations. There are no efficient general-purpose notations.

Researchers of language generation are in an advantageous position to investigate the adequacies and properties of various representation notation systems. This is so because English is itself organized around elementary varieties of knowledge that are highly recurrent, that are important to people, and are crucial in the solution of a great diversity of problems of everyday existence. English is a highly evolved notation, specialized over centuries of use to carry great varieties of knowledge. By developing and testing notations for linguistically prominent kinds of knowledge, researchers in language generation can help the development of solid and powerful knowledge representation systems.

During the course of the contract, we concentrated on three goals in this regard:

1. to expand existing knowledge notations to represent knowledge which is particularly crucial to generation;
2. to create new specialized notations to represent particular sorts of knowledge;
3. to develop techniques for reconciling and relating notations.

This work is described in Section 2.3.

2 Status of the Research Effort

2.1 Text Structure

The earliest feasible computational approach to the problem of producing coherent multi-sentence paragraphs involved the use of paragraph-sized structures called schemas which were essentially templates into which sentences could be fitted [McKeown 82]. Though effective and simple to describe and use, schemas suffer from some of the same limitations as templates do: they are not very flexible, and do not contain the rationale for the inclusion or order of any of their parts (which makes the parts themselves non-interchangeable).

Beginning in 1983 Dr. William Mann, then leader of the Penman project and PI on this contract, began developing a more suitable body of theory in collaboration with Prof. Sandra Thompson from the Linguistics Department at the University of California at Santa Barbara. This theory, called *Rhetorical Structure Theory* (RST) [Mann & Thompson 83], [Mann & Thompson 87a, Mann & Thompson 87b, Mann & Thompson 88a], [Mann & Thompson 88b, Mann 88a, Mann 88b, Mann 88c, Mann, Matthiessen & Thompson 88], [Thompson & Mann 86, Thompson & Mann 87] is based on the recognition that coherent English text exhibits an

internal structure. Any text can be broken into two or more principal parts, which themselves can each be broken down further, and so on recursively down to the single clause level. Adjacent clauses and blocks of clauses are related by rhetorical relations, of which (the claim is) English contains about 25. Thus, in RST texts are described by a dependency tree of rhetorical schemas, each of which relates several spans of text. Schemas are defined in terms of relations between a focal span, called the Nucleus, and Satellite spans. For example, one schema describes a span consisting of two smaller spans, one of which identifies a problem and the other of which identifies a solution to that problem. This schema is called *SOLUTIONHOOD*. Another schema, *EVIDENCE*, describes the combination of a claim and evidence for it. About 25 schemas have been identified, after a study of over 200 texts, spanning scientific article abstracts, cookbooks, letters, magazine articles, and so forth.

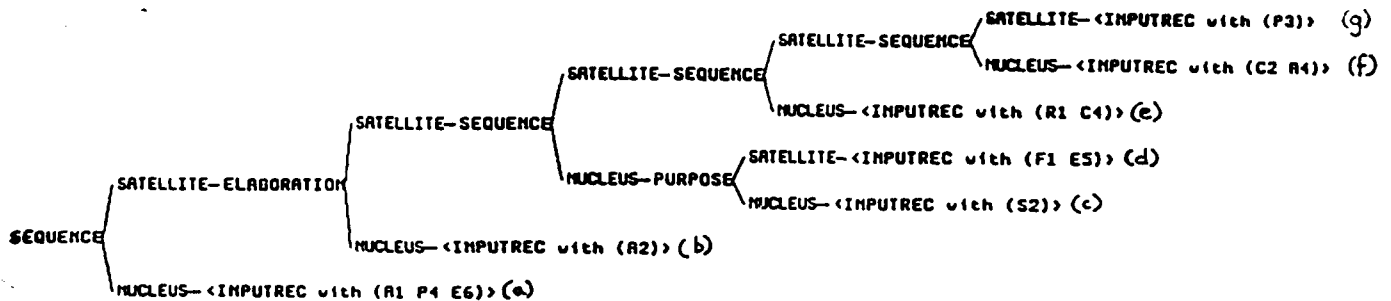
Since RST is recursive, it is capable of describing a wide range of sizes of text unit, from a single paragraph to a multiparagraph text such as a business report or magazine article. Whether the theory scales up to book-length texts is a matter for future study. On the small scale, though, the theory promises to help account for coherence properties, certain facts about text order, some aspects of thematization, tense, aspect, and many observations about conjunctions and pronominalization.

The year 1987 saw the first computational implementation of planning texts using RST [Hovy 88a, Hovy 88b]. In the implementation, the definitions of RST schemas used concept recognition criteria and representations of the speakers' goals, providing a basis for relating the coherence of text, as governed by RST relations, to the speaker's intentions in producing the text. As implemented, RST is essentially a goal-based theory. Its descriptions are organized around the intentions of the speaker and the part-to-part relations in the text which are used to carry out those intentions. In this it is comparable to recent work by Grosz and Sidner, but it does not work exclusively with the kind of fine-grain axiomatization of intention which they hope for [Grosz & Sidner 86].

The implemented planner is called the *text structurer* and has been tied successfully to the language generation program Penman. The structurer operates in top-down hierarchic expansion fashion, modeled on the planning system NOAH [Sacerdoti 77]. Its output is a tree that represents the internal dependencies of the parts of the paragraph, each of which is a piece of the input to the system. For example, as shown in Figure 1, an expert system that suggests changes to computer code to improve its readability and maintainability provides the planner with a collection of 7 units of information, gathered from its procedural knowledge, as well as the goal to explain the reasoning behind its recommendations. Using this goal to start planning, the text planner uses its library of RST relation/plans to build a tree in which branch points are RST relation/plans and leaves are input elements. It then traverses the tree, sending the leaves' content to the generator to be transformed into English. The tree in Figure 1 gives rise to the paragraph shown below the tree. It contains the RST relations *SEQUENCE* (signalled by "then" and "finally" in the paragraph), *ELABORATION* ("in particular"), and *PURPOSE* ("in order to").

The operationalization of the RST relations as plans is currently incomplete, both in

Figure 1: Paragraph structure tree for PEA text.



[The system asks the user to tell it the characteristic of the program to be enhanced.](a) Then [the system applies transformations to the program.](b) In particular, [the system scans the program](c) in order to [find opportunities to apply transformations to the program.](d) Then [the system resolves conflicts.](e) [It confirms the enhancement with the user.](f) Finally, [it performs the enhancement.](g)

the number of relations handled and in the combination of relation/plans with schemas for enhanced planning capability. We have operationalized only six of the twenty most basic RST relations. Operationalization involves formalizing the restrictions on a relation's use and the requirements for its parts in a language built from the formal theory of rational interaction currently being developed by, among others, Cohen, Levesque, and Perrault. For example, in [Cohen & Levesque 85], Cohen and Levesque present a demonstration that under certain assumptions the indirect speech act of requesting can be derived (recognized) using the following basic modal operators

- (BEL x p) — p follows from x 's beliefs
- (BMB x y p) — p follows from x 's beliefs about what x and y mutually believe
- (GOAL x p) — p follows from x 's goals
- (AFTER a p) — p is true in all courses of events after action a

We are using these relation/plans as compilations of these operators and the logical operations **AND** and **OR**. The operationalization task is difficult because one must ensure that the restrictions and requirements are formalized in ways that are at once specific enough to be directly useful in a computer program while being general enough to be applicable to the wide range of purposes for which the relations were originally intended.

The relation/plans we have formalized thus far — **SEQUENCE**, **ELABORATION**, **PURPOSE**, etc. — have enabled us to produce a number of paragraphs in three different domains of application (discussed in Sections 2.4.1, 2.4.2, and 2.4.3). Though these six relations have not been sufficient to produce all the kinds of texts one could produce from these domains, this method of planning coherent paragraphs has aroused considerable interest in the Natural Language Processing community. Various similar planners using RST and other relations have been built, both in the US and in Europe, and the nature, utility, and constraints of this technique are slowly beginning to become clear. Much work has been done by Moore and Paris at ISI in building their own text planner, initially using RST relation/plans, and then augmenting them with more semantically oriented domain-specific plans [Moore & Swartout 88, Moore & Swartout 89, Moore & Paris 89, Paris 88].

We have discovered that it is possible to formulate these plans as schemas, or even to form hybrids that are a mixture of schemas and plans. This finding is very encouraging, because it makes possible structuring tasks that otherwise are very difficult or impossible to perform. That is to say, since relation/plans are useful primarily when a large amount of flexibility is desired over a relatively small number (in the order of 10 to 30) of clause-sized units of information to be conveyed, they are less useful when faced with large collections of information. Here a less flexible method with more internal structure is required, if planning time is to be kept manageable — and this is exactly the strength of schemas. As explained in [Hovy 88b], it is possible to treat the growth points in relation/plans — those points that suggest the inclusion of additional material to the planning process — either as suggestions (in which case you get flexible planning) or as injunctions (in which case you get schemas).

This finding has not been extensively tested yet, and we have not integrated this notion into the planning system.

Another issue in which a preliminary investigation has been performed is the question of thematization or focus. Work done by McCoy and her students at the University of Delaware [McCoy & Cheng 88] suggested that the flow of focus or theme in coherent texts can be usefully represented in a hierarchical tree the called the Focus Tree. However, since Focus Trees do not take the rhetorical organization of paragraphs into account, they were not specific enough to support focus control alone. And since RST does not take the phenomena of thematization into account, as it has been developed to this point, a natural question arose about the possible conjoining of the two theories. A joint planning system was designed and described in [Hovy & McCoy 89], but further work awaits funding.

A number of other questions about text planning with RST and similar relation/plans and schemas remain to be addressed. These questions are summarized in [Hovy 89].

2.2 Parsing

Over the past three years, as part of this contract, a member of the Penman project has adapted a parser, built as part of his dissertation work [Kasper 87a, Kasper 87b], to use Penman's grammar Nigel [Kasper 88]. The first version of this parser was built by extending PATR-II [Shieber 84], a general unification-based system. It demonstrated that a general parsing capability could be developed for Systemic grammars that are expressed in a declarative notation.

However, the inefficiency inherent in the process of unification (the central process underlying the parser's action), a different approach was sought for. The formal similarity of the process of unification to the process of classification, as used in KL-ONE-like knowledge representation languages, suggested that if the grammar were to be represented in a such a representation system, unification could be replaced by classification, a much more efficient process. In such a scheme, the semantic and syntactic knowledge would be represented in the same system, accessed by the same process (unification), and potentially lead to a great simplification of parsing work in general. To perform this experiment, the representation system Loom developed at ISI [MacGregor & Bates 87] was chosen, it being on the forefront of KL-ONE-derived representation systems.

This work is still in progress; its completion awaits the incorporation of disjunction into the knowledge representation system Loom. When complete, both semantic information (as captured in the Upper and Domain Models; see below) and syntactic information (the grammar Nigel, as represented in Loom) will be accessible by the parser in a straightforward and homogeneous way. Furthermore, to aid the parsing process, Loom's classifier will replace the unification mechanism used previously by the parser.

The ability to perform inference over disjunctions is a necessary step in the parsing process. Until recently, inferences could not be made over disjunction (the logical operator

or) in the KL-ONE representation language family. This meant that grammars (as well as the intermediate structures built by parsers) could not be handled in these languages, because parsers necessarily deal with multiple options due to the structural and semantic ambiguities inherent in language. This inability was a serious problem, since the KL-ONE family of languages provide some of the best and most well-defined representation languages available. On the other hand, semantic knowledge is usually represented in these languages. The inability to represent both syntactic and semantic knowledge in the same system has precluded the development of parsers using a single inferencing technique such as classification to perform their work in a homogeneous and unified manner. Thus the lack of a general framework for computing with disjunctive knowledge structures has always been a hindrance to the development of parsing technology. Work is currently under way to incorporate inference over disjunctions into Loom at ISI (see [Kasper 87b, Kasper 88]), and this breakthrough will finally enable the representation of syntactic and semantic knowledge in the same representation system. This will enable parsers to access semantic and syntactic information as soon as it is relevant in a straightforward direct fashion using a single mechanism, the classifier, as a fast and efficient inferencing operation.

Though this work has not been completed, enough progress has been made to provide answers for some of the questions listed above, namely,

1. Bidirectionality: Could the inverse of the generation grammar be found?
2. Efficiency: Unconstrained Unification has a reputation for being exponentially slow in principle and extremely slow in practice. Would this be true of an inverted grammar?
3. Grammatical Specificity: No Systemic grammar had ever been examined for ambiguity behavior. Would we find large factors of preventable ambiguity?
4. Inversion-specific Inefficiencies: Does analysis using a grammar designed only for generation have unsuspected deficiencies in available information?

As a result of the research conducted under this project, we believe that a bidirectional systemic grammar must have a few small parts dedicated solely to generation and to analysis, but that nearly all of its parts can be shared. Thus the efforts of extending and maintaining this sort of bidirectional grammar are about the same as for a single-directional grammar.

The experiment of integrating a parser and generator with both semantic and syntactic knowledge represented in a KL-ONE-like representation system has never been carried out before. The imminent development of this capability is an exciting new breakthrough.

2.3 Knowledge Representation

To guide our work, we used the following criteria of importance and readiness:

1. Prefer varieties of knowledge that are required in every sentence over those that are optional.
2. Prefer varieties of knowledge whose representation will support other subtasks of this and related research.
3. Prefer varieties of knowledge for which the corresponding parts of our particular grammar are well elaborated.
4. Prefer varieties of knowledge for which we and/or others have already developed attractive proposals on how to represent the knowledge.

Based on these criteria, we concentrated on developing the representation of actions, their participants, and propositional ions.

As our basic knowledge representation system, we used two descendants of the KL-ONE knowledge representation formalism, namely NIKL and Loom. Both these systems were designed and built at ISI. Loom, a successor to NIKL which is still under construction, has many extensions over NIKL, which in turn had some desirable properties beyond those of KL-ONE. (The classifier, a mechanism which automatically classifies a newly-defined entity in terms of the existing definitions based on the aspects and properties of the entity, is one such feature.)

Our primary uses of the NIKL (and now, Loom) systems are to represent the information to be expressed. This information is typically produced by some application system, such as an expert system or a data base access program, which has to communicate its output needs to the generator using a language of common terms and structure.

The inputs given to a generator must be intelligible to it. Therefore, they must either be generator-internal symbols, or they must be defined in terms of symbols familiar and interpretable to the generator. To aid the user, Penman's decisions are formulated in terms of a taxonomy of the types of entities that appear in the world called the *Penman Upper Model*. The categories of the Upper Model reflect grammatical distinctions made in English (for example, actions are typically expressed as verbs and objects as nouns). Without such a taxonomy, Penman would have no way of determining whether to treat a user-defined symbol as an object, an action, a relation, etc.

In order to make use of domain-specific terms, the user must construct a *Domain Model* and subordinate it to the Upper Model. Upper Model entities define a very abstract partitioning of the world, and Domain Model entities define increasingly specific, more everyday, task-oriented distinctions. Together, the Upper and Domain Models constitute a generalization hierarchy organized as a property-inheritance network. When a new symbol is defined, it is placed in the taxonomy relative to one or more existing symbols, from whom it inherits features in addition to the particular features it is defined to have. Thus the user can formulate input to Penman using domain-specific terms, which, though themselves uninterpretable to Penman, inherit the features from the Upper Model that enable Penman to generate appropriate output.

The Upper Model: The top node of the Upper Model is the entity **THING**. The next level contains three subclasses, **OBJECT**, **QUALITY**, and **PROCESS** (which respectively organize objects (such as "ship"), qualities (such as "red", "operational"), and processes. The category **PROCESS** is divided into four types: **MENTAL-PROCESS**, **VERBAL-PROCESS**, **MATERIAL-PROCESS**, and **RELATIONAL-PROCESS**. Mental processes are such actions and states as "think" and "believe"; verbal processes are such actions as "tell" and "read"; material processes are the remaining actions and events, such as "sail" and "eat"; relational processes represent static relations, such as ownership, times, locations, etc.

By virtue of their positions in the inheritance hierarchy, entities inherit aspects or roles from their ancestors. Some of the commonly used aspects are:

- **domain:** any 2-place relation that is defined in the Upper or Domain Models holds between a domain and a range. This is the generic first place of a relation.
- **range:** The generic second place of a relation.
- **actor:** expresses the doer — the agent — of any **MATERIAL-PROCESS**, that is, of any action or event.
- **actee:** expresses the direct patient of any **MATERIAL-PROCESS**.
- **class-ascription:** expresses class membership (that is, the basic IS-A or A-KIND-OF subsumption relation).
- **property-ascription:** a general relation to express some property of an object.

The Upper Model is distinctive in that it reflects general category distinctions found in language in a way that most organizations of knowledge do not. For example, qualities are distinguished from classes (comparable to the adjective/noun distinction) rather than simply being treated alike as predicates. The Upper Model currently contains about 200 entities. With it, our system spans a large subset of English; in the three domains of experimentation we have tried it, we have not had to expand the Upper Model to any large extent.

The Domain Model: The Domain Model contains the definitions of the entities particular to the current application domain. This model should constitute a full ontology for the domain, defining all the types of objects, actions, relations, states, etc., that are used. Most applications require such a model as a natural part of their work, either explicitly or implicitly (for example, the field types and the relations among them in relational data bases).

These entity definitions in the domain model must be subordinated to entities in the Upper Model. That is to say, the entities defined in the Domain Model must form a hierarchy that can be knitted to the Upper Model in such a way that the inheritance proceeds smoothly down from Upper Model entities to increasingly specific Domain Model entities. Subordination provides the generation system with the general type of each domain entity

used. In addition, subordination provides the inheritance of the aspects (roles) that domain entities' ancestors take, as well as the accompanying constraints (number constraints, filler requirements, etc.). For example, if the entity CAR is subordinated to the entity VEHICLE, and VEHICLE is defined with the aspect AGE whose filler requirement is NUMBER, then a CAR will inherit the requirement that it have a numerical age. If the generator is then able to express the AGE aspect for one entity, it can express it for all entities. When a new entity is defined — say, FORD — and subordinated to CAR, then it inherits the AGE aspect and its filler requirement, and to be handled by Penman.

This inheritance of aspects and requirements is very useful. All the basic aspects of actions and relations have been defined in the Upper Model, which means that the user of the system has little additional work to do. Since every entity in the application domain must have (an) Upper Model ancestor(s), every domain entity will inherit a set of aspects from the Upper Model. (Of course, the entity may have additional domain-specific aspects as well.) This is one example of the power gained by a felicitous choice and use of knowledge representation system. We have developed an auxiliary program called UPPERMOST to facilitate the construction of the Domain Model.

These accomplishments serve our needs on action and participant relations particularly well, because they test action- and participant-oriented notations in both relatively language-neutral and relatively language-intense contexts. Recent project activity has involved coordinating these notations with other notions strongly related to action, such as events, times, places, outcomes, products and beneficiaries. They also serve well to test our ideas about propositional relations. Clause coordination in English and propositional relations in knowledge notation are in some ways two sides of the same coin. We have already been able to demonstrate relative clauses in English, along with English expression of time relations; we have also demonstrated several varieties of clause coordination. (Our generator's grammar currently allows 59 different kinds of clause combination, including 16 varieties of relative clause.)

It is important to note that the knowledge representation problem here is not a problem of whether the notations will in principle provide expressibility of particular information. Rather it is a problem of providing usable, manageable, compatible techniques for expressing a diversity of information.

The subsections below describe our approach to developing useful notations for particular varieties of knowledge.

2.3.1 Knowledge of Actions and Participants

English and other languages have elaborate provisions for describing actions and their participants. Of the two principal sentence types (Relational and Material), one is organized around an action expressed in the main verb, usually with other parts of the sentence devoted to identifying the participants in the action, such as its agent and the objects acted

upon. To be able generate texts, it is important to have control of the grammar of actions, and equally important to be able to represent efficiently the knowledge of actions.

AI's weakness in action representation is well recognized. One style of knowledge representation, the so-called frame oriented languages, are relatively well suited for the representation of action. However, it typically shows no strong differentiation between participant identification and other knowledge, and does not treat actions in a way that distinguishes them notationally from objects, states, relations, or other entities. The organization of natural languages suggests that there is a strong advantage to making such distinctions highly accessible. For example, many English words represent things from the point of view of participants in actions — words like "pilot", "researcher", and "observer"; there are also specialized suffixes used only to indicate participant roles: *grantor* and *grantee*. These enable communication and inference about actions, such as granting, independent of possible type distinctions (e.g., persons vs. institutions) among various participants. Rather than develop supplementary notations, we have extended an existing frame-oriented notation to provide more specifically for actions and their participants.

The use of the Upper Model is based on a strategy in which grammatical decisions are converted into taxonomic discriminations. Experience with this approach has been successful, but has also identified some problems. One of these problems arises because taxonomic distinctions derive from two sources: the linguistic conventions of English and the knowledge representation conventions of the host system for which sentences are generated. For example, a data base about travel may represent several kinds of trips: long and short trips, convention and conference trips, sales and recruiting trips. All of these may be represented in the data base as undifferentiated attributes of trips. Linguistically, *long* and *short* are attributes, best represented in the upper structure as qualities. Conventions and conferences are best represented as events, and sales and recruiting are best represented as kinds of processes or activities. Knowing these distinctions is essential to making the grammatical choices involved in talking about them.

The difficulty is that the regularity and homogeneity of the host system's knowledge needs to be retained, to keep it well organized and maintainable, but at the same time the linguistic differences need to be represented taxonomically for language generation. We are currently exploring several proposed solutions to this problem, but it has not yet been solved.

2.3.2 Knowledge of Propositional Relations

The expressive resources of English devoted to actions are strongly related to those devoted to propositions — roughly the expression of notions which take truth values. These resources are rich, including many methods for relating one proposition or action to another. The conjunctions and subordinators (including "and", "but", "when", "if", "although", "for instance", "that is", "so", "because", "then", "until", "while", and many more) are one part of this resource.

The weakness of AI notations in this area is well known. Notations oriented toward logic often do well with *and* and *or*, but the formal notation departs strongly from ordinary English usage. The other terms are more problematic. More diversity appears in the corresponding parts of frame-oriented notations, but there is relatively little language-oriented experience.

We approached this problem as follows:

1. In the Upper Model mentioned above, relations are given a distinguished place.
2. Within the relations subhierarchy, relations between propositions are given a distinguished place, and are further subdivided.
3. A small number of expressive facilities of the grammar are programmed to recognize particular interpropositional subtypes and employ the corresponding special facilities of English to express them.

The general strategy is to recognize in the high level knowledge organization conceptual distinctions that are important in English expression, and to use those distinctions in delivering the knowledge.

2.4 Testing the System: Collaborations within ISI

There is a methodological problem in developing knowledge delivery techniques: the techniques must somehow be tested and refined so that they work. Proofs of sufficiency-in-principle are not enough. The complexity of the subject makes it necessary to develop techniques on particular subject matter and knowledge rather than always working directly on the general case.

To meet this need we have begun to create a series of experimental text generation systems that embody the notations and processes being studied. The first of our series of experimental systems contained knowledge about computer versions of personal mail and appointment calendars. It was developed in conjunction with a related DARPA project which attempted to apply existing state-of-the-art technology to the problem of interacting with data bases in English. The DARPA project, part of the Strategic Computing program, served as a testbed for many of the ideas from Knowledge Delivery Research, and made it much easier to refine and extend these ideas.

With the development of an implementation of a planner that used some relations of RST, we had limited multisentential capability about two years ago. The mail and calendar system was replaced by collaboration with the following three projects (funded separately):

1. An integrated multimedia interface system (II), in which paragraphs of English text, planned and generated by Penman, are combined with maps, menus and other display methods, so as to be suitable for command and control use. As part of this work, a naval briefing environment was captured in which the English presented information

derived directly from a (sanitized) US Navy assets database. The project team was led by Dr. Norman Sondheimer.

2. The Program Enhancement Advisor (PEA) is an experimental expert system that interactively advises programmers on how their Lisp programs might be improved. It contains an explanation facility that uses Penman's grammar to generate text that explains how PEA works. PEA is being developed as a Ph.D. project by Johanna Moore under the direction of Dr. Bill Swartout.
3. The Digital Circuit Diagnosis system (DCD) is an experimental expert system that diagnoses faults in digital hardware. Like PEA, it contains an explanation facility that uses Penman's grammar to generate output. Text is generated that explains the definitions of entities within DCD and the reasoning that lead to the diagnosis. DCD is being developed by Dr. Cécile Paris in collaboration with Dr. Bill Swartout.

2.4.1 Application to Briefings from a Military Data Base

The first test of the multisentence planning capability was performed on data provided by the Integrated Interfaces application domain. In response to a user's request for information from a data base of Naval deployments, the II system gathered appropriate information and distributed it to its various output modes, one of which was the text planner and Penman. Some sample paragraphs generated by Penman in this domain were:

Knox, which is C4, is en route to San Diego in order to rendezvous with Task Group CTG70.1. It will arrive 4/24. It will perform exercises for four days.

Kennedy and Merrill are on a multisail to Sasebo, arriving 10/19. While it is in Sasebo, Kennedy, which is C4, will load until 10/22. Merrill will depart on 10/20 to be on operations until 10/30.

MEKAR-87 takes place in South China Sea from 10/20 until 11/13. Preble, Fanning, and Whipple are participating. Preble and Fanning arrive 10/20. Whipple arrives on 10/29. Preble, which is C3, will leave on 10/31. Fanning and Whipple will leave on 11/13.

2.4.2 Application to a Program Enhancement Advisory Tool

Beginning late in 1987, Penman was interfaced to the Program Enhancement Advisor (PEA), part of an independent research project at ISI. This step was particularly significant because PEA is a member of a design family of systems that are specially organized for

knowledge delivery. It is built in the Explainable Expert Systems framework, a generalization of foundational work by Dr. Bill Swartout.

It is commonly acknowledged that expert systems should be able to explain their behavior and methods, but most actual expert systems do so poorly, if at all. In the EES framework, programs are developed from the very beginning with explanation in mind, and much of the design information for a program is retained within it for use in explanations. After some initial use of our text structure planner, the PEA and DCD project members built their own text planner in roughly the same mold, affording them greater freedom of experiment, but continued to use the sentence generator Penman.

In addition to Figure 1, two texts from the PEA domain, describing the expert system's internal rules and process representations are the following (the structure of these paragraphs was planned by a text planner built by Moore and Paris [Moore & Paris 89]):

A transformation that enhances the readability of the program is defined as a transformation whose right hand side is more readable than its left hand side. One kind of a transformation whose right hand side is more readable than its left hand side is a transformation that has a right hand side that is a function that has a function name that is a common English word and a left hand side that is a function that has a function name that is a technical word. CAR-TO-FIRST is a transformation that has a right hand side that is a function that has a function name that is a common English word and a left hand side that is a function that has a function name that is a technical word.

The system asks the user to tell it the characteristic of the program to be enhanced. Then the system applies transformations to the program. In particular, the system scans the program in order to find opportunities to apply transformations to the program. Then the system resolves conflicts. It confirms the enhancement with the user. Finally, it performs the enhancement.

2.4.3 Application to a Digital Circuit Diagnosis System

In April, 1988 Penman was interfaced to a second program in the EES family, the Digital Circuit Diagnosis system (DCD), being developed by Dr. Cécile Paris. The DCD texts generated so far are definitional, and thus rely on different expressive facilities than PEA does.

Some interesting research has been conducted by Drs. Bateman (from the Penman project) and Paris (from the DCD project) on the generation of different surface forms of the same underlying propositional content. For example, the following three texts from the same underlying knowledge structure in the DCD domain, tailored to readers of various levels of sophistication, were generated by Penman:

The system is faulty, if there exists an O in the set of output terminals of the system such that the expected value of O does not equal the actual value of the signal part of O and for all I in the set of input terminals of the system, the expected value of the signal part of I equals the actual value of the signal part of I .

The system is faulty, if all of the expected values of its input terminals equal their actual values and the expected value of one of its output terminals does not equal its actual value.

The system is faulty, if the inputs are fine and the output is wrong.

The work of interfacing Penman to DCD was closely monitored so that we could understand the interfacing process. This led to a report which showed that interfacing currently takes about three person-weeks of effort, eventually reducible to about one person-week. Out of the experience of these two applications, and also to overcome some of Penman's internal notational problems, a new sentence specification language called SPL (Sentence Plan Language) was developed, to serve both as an internal notation between Penman's text planner and its sentence generator, and also as an external interface language for sentence generation. SPL has greatly reduced the amount of time it takes an outside user to learn to use Penman.

2.5 Collaborations outside ISI

In addition to collaborating with other projects within ISI, the Penman project is committed to getting Penman out to the community, both in order to have it used and tested, and in order to have other people work on extending the grammar. Currently, the Penman system runs in Common Lisp on Symbolics and TI Explorer Lisp machines, as well as on Macintosh-IIs and Sun workstations. It has been distributed to over 20 research sites throughout the world, and has been used as a focus of class instruction in graduate courses at Columbia University, the University of Delaware, and USC.

Recently, in order to promote increased development of various computational aspects of Systemic Linguistics, the project entered into a multinational collaboration in which various partners would have different focuses of research, while using Penman as a common

center. All work will be shared among all the partners and periodic updates will ensure that everyone is using the same basic mechanisms in their investigations. This collaboration, initiated by Dr. Erich Steiner, started in September 1989. The partners are:

- A group in the Linguistics Department of the University of Sydney, Australia
- The KOMET project at IPSI, Darmstadt, West Germany
- The Penman project at ISI, Los Angeles, USA

Roughly speaking, ISI will act as a clearing-house for the computational implementation and distribution of Penman and the parser, and support various aspects of research. IPSI will support research on generation and parsing as well. The Linguistics Department group in Sydney will pursue fundamental work on linguistic theory and grammar development. For more information, see Section 6.

2.6 New Opportunities in Knowledge Delivery

The success of the Penman system in planning and synthesizing texts opens up technical possibilities that were not previously available. In addition to knowledge delivery by means of synthesis of written English monologues, there are related communicative processes which might depend primarily on the same kinds of knowledge. These include synthesis of spoken English output, communication within interactive dialogue (especially in online human-computer interfaces) and various radical revisions of the underlying technology.

Each such change involves technical constraints on the methods used to achieve communication. Many of these constraints are unknown, so it is not clear what communicative possibilities are currently feasible.

It is now timely to explore some of these possibilities. We have identified several below. For each of these we expect to devote a small amount of effort to investigating the technical feasibility of extending present and forthcoming work in the given direction.

- **Speech Synthesis from Meanings:** The current capability for written text synthesis from meanings actually produces, as a by-product, much of the information that is needed for speech synthesis.
- **Dialogue and Interface Participation:** Engaging in dialogue or English-language human-computer interaction involves keeping track of a richer diversity of information about the other participant, and also a richer notion of communication planning, than monologue requires.
- **Multiple Perspectives:** One of the limitations of the techniques embodied in Penman is that there is a single fixed point of view toward each object in the system's knowledge. The view is selected at implementation time. This makes it difficult to use grammatical options such as nominalization, e.g. to use the verb "synthesize" pr the

noun "synthesis" when referring to the same process, or to use "those cows" instead of "that herd" to refer to the same group. Knowledge representation techniques that overcome this limitation are needed.

- **Alternative Control Structures:** Text generation is a complex problem involving a wide diversity of knowledge sources. Penman's control structure is a simple pipeline that attempts to anticipate all of the combinations that are important and likely. More effective control structures based on blackboards, unification, object-oriented programming, opportunistic inference and other techniques should be considered, especially for implementation of the remodularized system.

2.7 Summaries of the Principal Research Components

2.7.1 Knowledge Representation

Penman's Upper Model implements a taxonomic strategy for representing the linguistic expressive possibilities for specific kinds of knowledge. The strategy seems generally successful, but ongoing experimentation with this structure is needed to determine whether the strategy will work on very large or diverse collections of knowledge, and whether it will work when there is another organization imposed on the same body of knowledge.

The taxonomic strategy is being extended to a wide range of propositional relations, partly derived from RST, in order to test its effectiveness in a different way.

2.7.2 Text Structure

Work on constructing texts must rest on a strong descriptive theory. We now have such a descriptive theory, RST, in place, and it is being accepted by many linguists as a significant advance over what was previously available. The partial implementation of RST is useful in providing a model of how the descriptive theory can be made constructive, but the texts created so far are not big enough, diverse enough or numerous enough to judge the success of the implemented theory. These limitations can be overcome only with substantial effort in constructing experimental bodies of knowledge which are rich enough so that several interesting texts can be constructed for a given purpose. In addition, there must be attention to non-structural aspects of text planning in addition to the RST-related aspects, so that the quality of generated texts can be suitably evaluated.

These needs for extension and testing, for both knowledge representation and text structure, will be central research activities for the project in future research.

2.7.3 Parsing

The ability to handle inference over disjunctions in KL-ONE-like representation languages will have two major effects: greatly simplified parsers and enhanced processing speed and

efficiency. Development is in progress and is expected to be completed by the end of 1990. This innovation makes possible in the same KL-ONE-like representation systems the representation of both semantic knowledge (about application domains) and natural language grammars. In such schemes, the automatic concept classifier will be used as a powerful resource similar to the unifier to perform simultaneous syntactic and semantic-based classificatory inference under control of the parser. Until now, the flow of control between syntactic and semantic processing has always been a vexing question for parsers: for semantic processing, they have used classificatory inference (of various kinds) and for syntactic processing, a variety of other methods (including unification). The difficulty of making the results of each type of processing available to the other as soon as possible (since the two types of processing are mutually dependent) has always meant that one or the other process is made to perform more work (in some cases significantly more) than necessary, especially in maintaining numerous interpretation alternatives. In addition, the integration of the results of the two types of process into a common representation has required additional processing. The new integrated approach enabled by the ability to handle inference over disjunction has never been tried before. It is expected to simplify the parsing process considerably (since there is then only one inference process and its results are represented in a single formalism) and to increase the speed and efficiency of parsers (since each type of processing can be performed as soon as possible and no additional work need be done).

This ability is an exciting new development in the field of parsing and has aroused considerable interest in the Computational Linguistics community.

3 Partial List of Recent Publications

In the past two years (since January 1988), the Penman project has had over 60 publications in refereed journals and conferences. The following recent publications were written about the work sponsored under this contract:

1. Bateman, J., Kasper, R., Steiner, E. and Schütz, J. Interfacing an English Text Generator with a German MT Analyzer. In *Proceedings of the Annual Meeting of the GLDV*, Springer-Verlag, 1989.
2. Bateman, J., Kasper, R., Schütz, J. and Steiner, E. A New View on the Process of Translation. To be presented at the Conference of the European Association for Computational Linguistics, Manchester, England, April, 1989.
3. Bateman, J. Conversation Generation — a theoretical watershed? In *New Developments in Systemic Linguistics*, Fawcett, R. and Young, D. (eds.), Volume 2, Frances Pinter (to appear).

4. Bateman, J. and Paris, C. Phrasing a Text in Terms a User can Understand. In *Proceedings of the International Joint Conference of AI (IJCAI)*, Detroit, MI, August 1989.
5. Hovy, E. Planning Coherent Multisentential Text. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, Buffalo, NY, June, 1988. Also available as USC/Information Sciences Institute Reprint RS-88-208.
6. Hovy, E. Approaches to the Planning of Coherent Text. In *Proceedings of the 4th International Generation Workshop*, Los Angeles, CA, July, 1988.
7. Hovy, E. On the Study of Text Planning and Realization. In *Proceedings of the AAAI Workshop on Text Planning and Generation*, AAAI, St. Paul, MN, August, 1988.
8. Hovy, E.H. Unresolved Issues in Paragraph Planning. Presented at the Second European Workshop on Language Generation, Edinburgh, 1989. To appear in a book of selected papers from the workshop.
9. Hovy, E.H. and McCoy, K.F. Focusing your RST: A step toward generating coherent multisentential text. In the *Proceedings of the 11th Cognitive Science Conference*, Ann Arbor, 1989 (667-674).
10. Hovy, E., McDonald, D. and Young, S. Current Issues in Natural Language Generation: An Overview of the AAAI Workshop on Text Planning and Generation. In *AI Magazine* 10(3), 1989. Also in *Proceedings of the AAAI Workshop on Text Planning and Generation*, AAAI, St. Paul, MN, August, 1988.
11. Kasper, R. A Unification Method for Disjunctive Feature Descriptions. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Palo Alto, CA, July, 1987. Also available as USC/Information Sciences Institute Reprint RS-87-187.
12. Kasper, R. Conditional Descriptions in Functional Unification Grammar. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, Buffalo, NY, June 1988. Also available as USC/Information Sciences Institute Research Report RR-87-191, November, 1987.
13. Kasper, R. Ambiguity in Systemic Grammar: Experience with a Computational Parser for English. Presented at the 15th *International Systemics Congress*, East Lansing, MI, August 1988.
14. Kasper, R. An Experimental Parser for Systemic Grammars. In *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary, August 1988. Also available as USC/Information Sciences Institute Reprint RS-88-212.

15. Kasper, R. Systemic Grammar and Functional Unification Grammar. In *Systemic Functional Approaches to Discourse*, Benson, J. and Greaves, W. (eds), Norwood, NJ: Ablex (in press). Also available as USC/Information Sciences Institute Reprint RS-87-189.
16. Mann, W. and Thompson, S. Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, Volume 1, 1987. Also available as USC/Information Sciences Institute report RS-87-185.
17. Mann, W. and Thompson, S. Rhetorical Structure Theory: Description and Construction of Text Structures. In *Natural Language Generation: New results in Artificial Intelligence, Psychology and Linguistics*, Kempen, G. (ed.), Nijhoff, 1987. Also available as USC/Information Sciences Institute report RS-87-174.
18. Mann, W. Text Generation: The Problem of Text Structure. In *Natural Language Generation Systems*, McDonald, D. and Bolc, L. (eds), Springer-Verlag: New York, 1988.
19. Mann, W. Dialogue Games. In *Argumentation*, 1988. Also available as USC/Information Sciences Institute report RR-79-77.
20. Mann, W. Two Theories of Discourse Structure. In *Proceedings of the 4th International Generation Workshop*, Los Angeles, CA, July, 1988.
21. Mann, W., Matthiessen, C. and Thompson, S. Rhetorical Structure Theory and Text Analysis. In *Discourse Description: Diverse Analyses of a Fund Raising Text*, Mann, W. and Thompson, S. (eds), (to appear).
22. Mann, W. and Thompson, S. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. In *Text*, Vol. 8:3, 1988.
23. Mann, W. and Thompson, S. Rhetorical Structure Theory: A Theory of Text Organization. In *The Structure of Discourse*, Polanyi, L. (ed), Ablex: Norwood, NJ, 1988.
24. Matthiessen, C. Lexical Selection in Generation: An Abstract Model. In *Proceedings of the 4th International Generation Workshop*, Los Angeles, CA, July, 1988.
25. Matthiessen, C. Representational Issues in Systemic Functional Grammar. In *Systemic Functional Approaches to Discourse*, Benson, J. and Greaves, W., Ablex, 1988. Also available as USC/Information Sciences Institute report RS-87-179.
26. Matthiessen, C. and Thompson, S. The Structure of Discourse and Subordination. To appear in *Clause Combining*, Haiman, J. and Thompson, S. (eds), Amsterdam: John Benjamins, 1988.

27. Matthiessen, C. Representational Issues in Systemic Functional Grammar. In *Systemic Functional Approaches to Discourse*, Benson, J. and Greaves W. (eds), Norwood, NJ: Ablex (in press).
28. Sondheimer, N., Cumming, S. and Albano, R. How to Realize a Concept: Lexical Selection and the Conceptual Network in Text Generation. In *Proceeding of the Workshop on Theoretical and Computational Issues in Lexical Semantics*. Waltham, Massachussets, April, 1988.
29. Sondheimer, N. Lexical Selection. In *Proceedings of the Workshop on Conceptual Networks*, AAAI, St. Paul, MN, August, 1988.
30. Thompson, S. and Mann, W. A Discourse View of Concession in Written English. In *Proceedings of the Second Annual Meeting of the Pacific Linguistics Conference*, DeLancey, S. and Tomlin, R., 1986.
31. Thompson, S. and Mann, W. Antithesis: A Study in Clause Combining and Discourse Structure. In *Language Topics: Essays in Honour of M.A.K. Halliday*, Steele, R. and Threadgold, T. (eds.), Benjamins, 1987. Also available as USC/Information Sciences Institute report RS-87-171.
32. Whitney, R. *Semantic Transformations for Natural Language Production*. USC/Information Sciences Institute report RR-88-192, March, 1988.

In addition to the publications listed in the previous section, the following presentations were made about work sponsored under this contract:

1. Mann, W. and Matthiessen, C. Functions of Language in Two Frameworks. Presented at the 14th *International Systemics Congress*, Sydney, Australia, July, 1987.
2. Matthiessen, C. and Mann, W. Rhetorical Structure Theory and Systemic Approaches to Text Generation. Presented at the 14th *International Systemics Congress*, Sydney, Australia, July, 1987.
3. Bateman, J., Kasper R. and Matthiessen, C.M.I.M. Systemic Linguistics and Natural Language Processing: Case Studies in the Exchange. Invited workshop presentation at the 15th *International Systemics Congress*, East Lansing, MI, August, 1988.
4. Bateman, J. Dynamic Systemic Functional Grammar: A New Frontier. Presented at the 15th *International Systemics Congress*, East Lansing, MI, August, 1988.
5. Kasper, R. Ambiguity in Systemic Grammar: Experience with a Computational Parser for English. Presented at the 15th *International Systemics Congress*, East Lansing, MI, August 1988.

6. Mann, W. Two Approaches to Discourse Structure from Computational Linguistics. Presented at the 15th *International Systemics Congress*, East Lansing, MI, August 1988.
7. Matthiessen, C. Notes on the Organization of the Environment of a Text Generation Grammar. In *Natural Language Generation: New results in Artificial Intelligence, Psychology and Linguistics*, Kempen, G. (ed.) Nijhoff, 1987. Also available as ISI report ISI/RS-86-177.
8. Hovy, E.H. Unresolved Issues in Paragraph Planning. Presented at the Second European Workshop on Language Generation, Edinburgh, 1989.
9. Hovy, E.H. and McCoy, K.F. Focusing your RST: A step toward generating coherent multisentential text. Poster presentation at the *11th Cognitive Science Conference*, Ann Arbor, 1989.

4 Personnel

The following personnel were supported in full or in part in the duration of this contract (degrees listed were attained under partial sponsorship of this contract; recipients were either part-time project members before graduation or joined the project full-time afterward):

- Mr. Robert N. Albano (currently graduate student at UCLA)
- Dr. John A. Bateman (currently project member, working in Germany)
- Ms. Susanna Cumming (currently a Linguistics Department faculty member at the University of Colorado; attained Ph.D. in Linguistics from UCLA in May 1987)
- Mr. Tom Y. Galloway (currently working in Geneva)
- Dr. Eduard H. Hovy (currently project leader; attained Ph.D. in Computer Science from Yale University in May 1987)
- Dr. Robert T. Kasper (currently project member; attained Ph.D. in Computer Science from the University of Michigan in December 1986)
- Ms. Lynn Poulton (currently a Linguistics Department graduate student at the University of Sydney)
- Dr. William C. Mann (on partial retirement)
- Mr. Christian M.I.M. Matthiessen (currently a Linguistics Department faculty member at the University of Sydney; attained Ph.D. in Linguistics from UCLA in December 1988)
- Dr. Norman K. Sondheimer (currently head of the AI division of GE Corporate Research)
- Mr. Richard A. Whitney (currently project member; attained M.S. from UCLA in May 1988)

5 Interactions and Meetings

The 4th International Workshop on Natural Language Generation was held in July, 1988, with USC as sponsor (along with AAAI, ACL and ACM), hosted by Drs. Mann, Paris and Swartout from ISI. Although no AFOSR funds were applied to workshop expenses, the project benefited from extensive interactions with leaders in this kind of work. After the workshop, most of the conferees visited ISI for 1 to 3 days, which included numerous demonstrations of Penman and other research systems.

Our past research has gained enormously from visiting workers who have had no formal status on the project. Several eminent and highly qualified people have visited for periods of weeks, without pay, relating their work and expertise to the ongoing research. Visitors who stayed for at least two weeks include, from The Federal Republic of Germany Drs. H-J. Novak, B. Nebel, E. Steiner, J. Schütz; from Britain Ms. J. Wright; from Yugoslavia Dr. M. Simunović. Other visitors included Drs. D. Rösner, G. Kempen, K. Sparck-Jones, D. Weber, K. Shimohara; Messrs. N. Reithinger and M. Elhadad; and Ms. C. DiMarco.

Two project members spent three months of 1989 working in Europe. Dr. Eduard Hovy was invited by the IBM Natural Language Research Laboratory in Stuttgart, West Germany, to continue work on text planning, and Dr. John Bateman was invited by EUROTRA-D, the German European machine translation project EUROTRA to work at their lab in Saarbrücken, West Germany. During this time both project members travelled widely and presented a number of talks at various institutions, including the Universities of Stuttgart, Saarland, Bielefeld (all three in Germany), and Linköping (Sweden).

6 Collaborations

In addition to using Penman within ISI, project members have been collaborating with various researchers from other institutions. In the past year, the following collaborations occurred:

- The Penman project entered into a collaboration with the Komet project at IPSI, a research institution in Darmstadt, West Germany, funded as part of the country-wide Institute for Mathematics and Computer Science research GMD by the federal German government. A Penman project member has spent the past five months working at IPSI on joint research, and a visitor from IPSI arrives at ISI later this month.
- The Penman project also entered into a collaboration with the Linguistics Department at the University of Sydney, Sydney, Australia, which is the home of Systemic-Functional Linguistics, the basis of our grammar. Linguists at Sydney collect and collate grammar development efforts by Systemic Linguists from around the world, and then pass the result on to Penman to be incorporated into the grammar.

- EUROTRA-D, the German branch of the European machine translation project, after sending three project members to visit ISI for 6 weeks, have decided to use Penman for English generation. EUROTRA is a five-year EEC-wide project whose aim is to produce systems that translate technical documents for EEC nations.
- A collaboration has been started with Prof. Kathy McCoy from the University of Delaware, whose theory of focus augments the research we have been performing on multisentence text planning.

7 System Distribution

The Penman text generation system has recently been structured as a distributable system, and has been distributed to over 15 institutions to date, including:

- Columbia University, New York City: being used in graduate seminar
- University of Toronto, Toronto: planned use in the thesis work of at least one graduate student
- University of Delaware, Newark: being used in graduate seminar
- Sydney University, Sydney, Australia: used for grammar development
- IPSI, Darmstadt, West Germany: language generation research
- EUROTRA-D, Saarbrücken, West Germany: awaiting porting to Sun computer
- University of Alabama, Huntsville: used in Ph.D. study of a student
- New Mexico State University, Las Cruces
- University of Berlin, Berlin, West Germany: graduate study
- University of Stuttgart, Stuttgart, West Germany: graduate study
- York University, Toronto: being ported to Vax computer

The following institutions have requested or expressed preliminary interest in Penman, but have not yet completed the licensing agreement:

- University of California, Berkeley.
- Carnegie-Mellon University, Pittsburgh.
- University of Massachusetts, Amherst.

The following institutions have, or have expressed the desire to acquire, a paper (non-computer-based) copy of Penman's grammar:

- IBM Natural Language Center, Los Angeles: currently installing the grammar on their own systems.
- York University, Toronto (English Department).
- University of Wales, Cardiff (English Department).

8 New Discoveries and Inventions

No new inventions or patent disclosures resulted from this work.

9 Other Statements Assisting Evaluation

No other statements are required to provide additional insight and information for an assessment of the work done under this contract.

10 References

References

- [Cohen & Levesque 85] Cohen, P.R. & Levesque, H.J. Speech Acts and Rationality. In *Proceedings of the ACL Conference*, Chicago, 1985.
- [Grosz & Sidner 86] Grosz, B.J. & Sidner, C.L. Attention, Intentions, and the Structure of Discourse. In *Computational Linguistics Journal* 12(3), 1986.
- [Hovy 88a] Hovy, E.H. Planning Coherent Multisentential Text. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, Buffalo, NY, June, 1988. Also available as USC/Information Sciences Institute Reprint RS-88-208.
- [Hovy 88b] Hovy, E.H. Approaches to the Planning of Coherent Text. In *Proceedings of the 4th International Generation Workshop*, Los Angeles, CA, July, 1988.
- [Hovy 89] Hovy, E.H. Unresolved Issues in Paragraph Planning. Presented at the Second European Workshop on Language Generation, Edinburgh, 1989. To appear in a book of selected papers from the workshop.

- [Hovy & McCoy 89] Hovy, E.H. and McCoy, K.F. Focusing your RST: A step toward generating coherent multisentential text. In the *Proceedings of the 11th Cognitive Science Conference*, Ann Arbor, 1989 (667-674).
- [Kasper 87a] Kasper, R.T. A Unification Method for Disjunctive Feature Descriptions. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Palo Alto, CA, July, 1987. Also available as USC/Information Sciences Institute Reprint RS-87-187.
- [Kasper 87b] Kasper, R.T. Conditional Descriptions in Functional Unification Grammar. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, Buffalo, NY, June 1988. Also available as USC/Information Sciences Institute Research Report RR-87-191, November, 1987.
- [Kasper 88] Kasper, R.T. An Experimental Parser for Systemic Grammars. In *Proceedings of the 12th International Conference on Computational Linguistics*, Budapest, Hungary, August 1988. Also available as USC/Information Sciences Institute Reprint RS-88-212.
- [MacGregor & Bates 87] MacGregor, R. and Bates, R. The Loom Knowledge Representation Language. In *Proceedings of the Knowledge-Based Systems Workshop*, St. Louis, 1987. Also available as USC/Information Sciences Institute Research Report RS-87-188, 1987.
- [Mann 88a] Mann, W.C. Text Generation: The Problem of Text Structure. In *Natural Language Generation Systems*, McDonald, D. and Bolc, L. (eds), Springer-Verlag: New York, 1988.
- [Mann 88b] Mann, W.C. Dialogue Games. In *Argumentation*, 1988. Also available as USC/Information Sciences Institute report RR-79-77.
- [Mann 88c] Mann, W.C. Two Theories of Discourse Structure. In *Proceedings of the 4th International Generation Workshop*, Los Angeles, CA, July, 1988.
- [Mann, Matthiessen & Thompson 88] Mann, W.C., Matthiessen, C.M.I.M. & Thompson, S.A. Rhetorical Structure Theory and Text Analysis. In *Discourse Description: Diverse Analyses of a Fund Raising Text*, Mann, W. and Thompson, S. (eds), (to appear).
- [Mann & Thompson 83] Mann, W.C. & Thompson, S.A. Relational Propositions in Discourse. USC/Information Sciences Institute Research Report RR-83-115, 1983.

- [Mann & Thompson 87a] Mann, W.C. & Thompson, S.A. Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, Volume 1, 1987. Also available as USC/Information Sciences Institute report RS-87-185.
- [Mann & Thompson 87b] Mann, W.C. & Thompson, S.A. Rhetorical Structure Theory: Description and Construction of Text Structures. In *Natural Language Generation: New results in Artificial Intelligence, Psychology and Linguistics*, Kempen, G. (ed.), Nijhoff, 1987. Also available as USC/Information Sciences Institute report RS-87-174.
- [Mann & Thompson 88a] Mann, W.C. & Thompson, S.A. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. In *Text*, Vol. 8:3, 1988.
- [Mann & Thompson 88b] Mann, W.C. & Thompson, S.A. Rhetorical Structure Theory: A Theory of Text Organization. In *The Structure of Discourse*. Polanyi, L. (ed), Ablex: Norwood, NJ, 1988.
- [McCoy & Cheng 88] McCoy, K.F. and Cheng, J. Focus of attention: Constraining what can be said next. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [McKeown 82] McKeown, K.R. *Generating Natural Language Text in Response to Questions about Database Queries*. Ph.D. dissertation, University of Pennsylvania, 1982.
- [Moore & Paris 89] Moore, J.D. and Paris, C.L. Planning Text for Advisory Dialogues. In *Proceedings of the 27th ACL Conference*, Vancouver, 1989.
- [Moore & Swartout 88] Moore, J.D. and Swartout, W.R. Dialogue-based explanation. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [Moore & Swartout 89] Moore, J.D. and Swartout, W.R. A reactive approach to explanation. To be presented at the *International Joint Conference on Text Generation IJCAI*, Detroit, 1989.
- [Paris 88] Paris, C.L. Generation and explanation: Building an explanation facility for the Explainable Expert Systems framework. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [Sacerdoti 77] Sacerdoti, E. *A Structure for Plans and Behavior*. North-Holland, 1977.

- [Shieber 84] Shieber, S.M. The design of a computer language for linguistic information. In *Proceedings of the Tenth International Conference on Computational Linguistics (COLING)*, Stanford, 1984.
- [Thompson & Mann 86] Thompson, S.A. & Mann, W.C. A Discourse View of Concession in Written English. In *Proceedings of the Second Annual Meeting of the Pacific Linguistics Conference*, DeLancey, S, and Tomlin, R., 1986.
- [Thompson & Mann 87] Thompson, S.A. & Mann, W.C. Antithesis: A Study in Clause Combining and Discourse Structure. In *Language Topics: Essays in Honour of M.A.K. Halliday*, Steele, R. and Threadgold, T. (eds.), Benjamins, 1987. Also available as USC/Information Sciences Institute report RS-87-171.